

-1-

**METHODS AND APPARATUS FOR DECODING IMAGES USING DEDICATED
HARDWARE CIRCUITRY AND A PROGRAMMABLE PROCESSOR**

FIELD OF THE INVENTION

The present invention relates to decoding of images and, more specifically, to methods and apparatus
5 for performing video decoding in a system including a programmable processor, e.g., a central processing unit (CPU).

BACKGROUND OF THE INVENTION

10 The digital representation of images is becoming ever more common. Digital video, e.g., digital high definition television (HDTV), represents one application where digital video signals comprise, e.g., a
15 sequence of digital images. With the increasing use of the Internet, the generation, storage and transmission of digital video images is likely to continue to increase.

In order to reduce the amount of data required
20 to represent digital images, such images are often encoded, e.g., compressed. Coding of images in the form of frames may be performed on an intra-frame basis so that the generated encoded image data does not depend on data from other images. Intra-frame coding is often
25 performed using discrete cosine transform (DCT) coding

techniques. Data generated from intra-frame coding is often called intra-coded data.

When a sequence of images is being coded,
5 coding of images may be performed on an inter-frame, as well as an intra-frame, basis. Inter-frame coding generally involves using motion compensated prediction techniques to produce motion vectors which include information on how a portion of a subsequent image may be
10 generated using a portion of a preceding or subsequent image as reference data. Motion vector information, once generated, may be coded, e.g., using differential coding techniques, to reduce the amount of data required to represent the motion vector information.

15 When there is little motion between frames, motion vectors provide an efficient method of representing portions of image. However, intra-frame coding can be more efficient when there is a substantial
20 change from one frame to another. In an attempt to maximize coding efficiency, often some portions of an inter-coded frame are encoded using intra-frame coding techniques while other portions of the same image are coded using inter-frame coding, e.g., motion vectors.
25 Those portions of inter-coded frames that use inter-frame coding, may also include coded information regarding the residual, or correction, image. During decoding of such portions of frames, the coded residual data is decoded. The decoded residual image data is combined with image

data generated by performing motion compensated predictions using reference frame data and motion vectors. Those portions of inter-coded frames that use intra-frame coding are decoded without the use of
5 reference frame data or motion vectors. In this manner, a complete inter-coded frame can be generated from the intra- and inter-coded data used to represent the frame.

5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63 64 65 66 67 68 69 70 71 72 73 74 75 76 77 78 79 80 81 82 83 84 85 86 87 88 89 90 91 92 93 94 95 96 97 98 99 100

MPEG-2 is a well known video standard which includes support for the use of motion compensated prediction techniques in addition to transform coding, e.g., the use of DCT transforms. The international video standard MPEG-2 is described in the International Standards Organization document (ISO/IEC 13818-2). MPEG-2 has been used as the basis for several commercial applications including digital video disks (DVD) and digital broadcast television.

In MPEG-2 motion vector information is
20 differentially encoded for transmission purposes. In addition, quantization operations and run length coding operations are performed on video data to further reduce the amount of data required to represent the images being encoded. A scan conversion operation is also normally
25 performed as part of the MPEG-2 coding process in order to convert a set of two dimensional coefficient data into a one dimensional data sequence which can be processed, stored and/or transmitted. Variable length coding is

applied to many of the data elements in order to further reduce the number of bits needed to represent an image.

In order to view images represented using
5 encoded video data, the image data has to be decoded
prior to display. MPEG-2 decoding generally involves
performing operations which are the inverse to those used
to originally encoded the image data, e.g., decoding
usually involves variable length decoding, inverse scan
10 conversion, inverse quantization, inverse discrete cosine
transform (IDCT), motion vector reconstruction and motion
compensated prediction operations.

Given the large amount of data used to
15 represent images, a considerable amount of processing is
normally required to decode an encoded image. In the
case of motion vectors, image decoding is complicated by
the need to access reference frame data, e.g., previously
decoded frames, in order to generate a current decoded
20 frame. The need to access reference frame data to
perform motion compensated predictions using motion
vector information results in the motion compensated
prediction operations being highly memory intensive. The
large number of memory access operations associated with
25 performing motion compensated predictions can often have
a significant impact on the amount of resources required
to decode an image.

Users of computers are beginning to expect that they will be able to decode and display video images, e.g., encoded motion pictures, in real time. In order to achieve the real time display of encoded video, the encoded video images, e.g., frames, on average, need to be decoded in the same or a smaller amount of time than is used to display the images. As discussed above, the amount of processing required to decode an image can be considerable. Placing real time restraints on decoder circuitry further complicates matters due to the time constraints in which the decoding must be accomplished.

Generally, known video decoders belong to one of three types: 1) those that use dedicated special purpose integrated decoder circuits coupled with video memory to fully decode encoded video data; 2) those that use software and a general purpose programmable processor, e.g., CPU such as a Pentium processor, to fully decode encoded video data; and 3) those that partition the video decoding between a general purpose programmable processing unit such as a Pentium processor and a graphics processor chip which is used to manage and/or perform memory intensive operations such as motion compensated prediction. A high speed video or graphics memory is often used in conjunction with the graphic's processor to further accelerate image decoding operations.

In many high definition television sets and other video display devices, the first type of known decoders are used. That is, dedicated hardware decoder circuits are used to decode received encoded video data.

5 Fig. 1 illustrates a known decoder circuit 100 capable of decoding MPEG-2 video signals.

As illustrated, the video decoder 100 comprises a variable length decoder (VLD) circuit 102, an inverse scan circuit 104, an inverse quantization circuit 106, an
10 inverse discrete cosine transform (IDCT) circuit 108, a motion compensated prediction circuit 110, and a frame store memory 112. Motion vector reconstruction circuitry, which is used to reconstruct motion vectors
15 from encoded motion vector information, is present but not explicitly shown in the Fig. 1 illustration.

In the known system, the VLD circuit 102 receives MPEG-2 encoded video data and performs a
20 variable length decoding operation on variable length encoded data included therein. The inverse scan circuit 104 is responsible for re-sequencing elements of the video data output by the VLD circuit 102 to reverse the effect of the scan conversion operation originally used
25 to convert the two dimensional coefficient data into a one dimensional ordering. The inverse quantization (IQUANT) circuit 106 performs an inverse quantization operation on the quantized values included in the video data output by the inverse scan circuit 104. The IQUANT

circuit 106 generates a video data stream including DCT coefficients. The DCT coefficients are processed by the IDCT circuit 108 to generate decoded image data from received intra-coded image data. The output of the IDCT circuit, in the case of intra-coded data, is decoded video, e.g., pixel values which are used to control the light output of individual pixels.

As discussed above inter-coded images may include both intra and inter coded image portions. The motion compensated prediction circuit is responsible for using previously decoded image data, which has been stored in the frame store memory 102, as reference frame data. The motion compensated prediction circuit 110 outputs received decoded image data corresponding to intra-coded frames as well as decoded image data corresponding to inter-coded frames which it generates from received motion vectors and/or by combining decoded inter-coded image data with decoded image data.

20

In the case of personal computers, the computer already includes a general purpose processor which is capable of being used for video decoding operations. Accordingly, existing personal computers usually use the second or third of the known video decoding techniques discussed above.

25

Unfortunately, software decoder implementations tend to be far slower than hardware implementations

operating at comparable clock speeds. This is because the CPU circuitry is not optimized for video decoding as is the case with a dedicated hardware circuit. In addition, the relatively large amount of data which must
5 be transferred from memory when performing motion compensated prediction operations tends to create a bottleneck with regard to decoding speed.

While the use of a graphics processor and a
10 dedicated video memory can help improve video decoding performance, the use of a general purpose CPU to perform non-memory intensive decoder operations still normally results in a slower decoder implementation than would be achieved using dedicated decoder hardware operating at
15 comparable clock speeds.

Given the processing capability of CPU's currently used in common personal computers, and the increasing tendency to run multiple applications at the
20 same time, real time software based video decoding of HDTV and other high resolution video images is often impractical on modern personal computers.

In view of the above, it becomes apparent that
25 there is a need for improved methods and apparatus for decoding video images in systems, e.g., personal computers, which include programmable CPUs. It is desirable that at least some of the new methods and/or apparatus be capable of being implemented at a lower cost

than providing a complete dedicated video decoder circuit capable of decoding intra-coded and inter-coded image data. It is also desirable that at some of the new methods and apparatus be capable of being used with
5 systems which include a graphics processor for performing and/or overseeing memory intensive decoding functions such as motion compensated prediction operations.

SUMMARY OF THE INVENTION

10

The present invention is directed to methods and apparatus for decoding images. In accordance with the present invention video decoder circuitry is combined with a programmable processor, e.g., a general purpose
15 CPU, to perform video decoding operations. A graphics processor may be used in combination with the CPU and video decoder circuitry of the present invention to perform and/or oversee memory intensive video decoder operations.

20

The video decoder circuitry of the present invention performs one or more non-memory intensive video decoding operations such as syntax parsing, variable length decoding, DC DCT coefficient reconstruction,
25 motion vector reconstruction, inverse scan operations, inverse quantization and/or mismatch control. Memory intensive video decoding operations such as performing prediction operations using motion vectors and reference frames are performed by the programmable processor and/or

graphics processor operating in conjunction with the video decoding hardware of the present invention. Because the functions performed by the video decoder circuitry of the present invention are normally performed
5 at the beginning portions of the inter-coded decoding process, e.g., prior to performing motion compensated predictions, the video decoder circuitry of the present invention may be characterized as a video decoder front end processor. In various embodiments, the video decoder
10 front end processor includes a complete intra-coded decoder. In addition it may, and in various embodiments does, include a motion vector reconstruction circuit.

The decoder circuitry of the present invention
15 does not perform memory intensive motion compensated predictions. Accordingly, it can be implemented far more cheaply than video decoder circuits which perform such operations.

20 By combining the relatively low cost video decoder circuitry of the present invention with a general purpose processor, superior decoded video quality can be achieved and/or more images decoded in a given amount of time than when a similar general purpose processor is
25 used to decode images without the benefit of the decoder circuitry of the present invention. Thus, the methods and apparatus of the present invention provide a way of enhancing a computer system's ability to decode video

images at a fraction of the cost of adding a hardware decoder capable of fully decoding inter-coded images.

5 The video decoder circuitry of the present invention can be implemented on a separate chip from the computer system's CPU. Such a chip can be incorporated into a computer card which can be inserted into a computer to enhance the system's ability to perform video decoding. Alternatively, the video decoder circuitry of
10 the present invention can be incorporated directly into a processor thereby providing the video decoder circuitry of the present invention and a programmable processor on a single semiconductor chip.

15 Various additional features and advantages of the present invention will be apparent from the detailed description which follows.

20 BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 illustrates a known MPEG-2 video decoder.

25 Figs. 2 and 3 illustrate exemplary computer systems of the present invention wherein a video front end processor is used in combination with a CPU and software to decode encoded video.

Fig. 4 illustrates the front end processor of the present invention, used in the Fig. 2 and 3 embodiments, in greater detail.

5

DETAILED DESCRIPTION

As discussed above, the present invention is directed to methods and apparatus for decoding images, e.g., MPEG-2 video images. The methods and apparatus of the present invention are well suited for use with systems which include a programmable processor, e.g., a computer system which includes a CPU. They are also well suited for use in systems which include, in addition to a CPU, a graphics processor for performing or overseeing memory intensive video decoding operations.

In contrast to prior art techniques which used a general purpose processor to perform non-memory intensive video decoding operations, the inventors of the present application realized that there can be processing speed and quality advantages to performing non-memory intensive video decoding operations in hardware. It was further recognized that a hardware circuit to perform such decoder operations could be implemented at relatively little cost and could be used in conjunction with a general purpose processor, or incorporated into a general purpose processor, to facilitate video decoding operations.

By using dedicated video decoder hardware circuitry to perform some, most, or all of the non-memory intensive operations associated with decoding one or more
5 images, the ability of a computer system to decode still images and video is enhanced at a fraction of the price of providing a full hardware video decoder circuit capable of decoding inter-coded and intra-coded images.

10 Figure 2 illustrates a computer system 300 implemented in accordance with one embodiment of the present invention. As illustrated the system 300 includes a video decoder front end processor circuit 302 used for performing non-memory intensive video decoding
15 operations in accordance with the present invention, a programmable processor (CPU) 316, a display device 210, display adapter 312, keyboard 234, a first memory bus 229, a second memory bus 229', and a system bus 228.

20 The system 300 further includes a secondary (L2) cache 236 which is coupled to the CPU 316 via the first memory bus 229 and to the main memory 326 via the second memory bus 229'. The system bus 228 couples a network adapter 224 and a peripheral device 322 to one
25 another and to the CPU 316. The network adapter 224 is used to couple the system 300 to a digital network via which encoded video data may be received.

The peripheral device 322 is coupled to the video front end processor circuit 302 in addition to being coupled to the system bus 228. In this manner, the peripheral device 322 which may be, e.g., a video disc player, can supply encoded video directly to the video front end processor circuit 302. While a single peripheral device 322 is shown in the system 300 it is to be understood that the system can, and in various embodiments does, include multiple peripheral devices 322.

In the system 300, the CPU 316 includes a cache 320 and an I/O interface 318. The I/O interface 318 couples the CPU 316 to the video front end processor circuit 302, display adapter 312, keyboard 234, system bus 228 and memory bus 229. The I/O interface 318 serves to electrically connect the devices coupled thereto and to convert between internal signaling formats used by the CPU 316 and external devices coupled thereto. The CPU 316 performs processing operations under control of routines stored in the main memory 326. The stored routines include a video decoder routine 330 and various other routines, e.g., spread sheet, word processor, web browser, etc. As will be discussed below, unlike the prior art video decoder routines which were responsible for controlling a CPU to fully decode inter-coded video data, the video decoder routine 330 of the present invention is designed to control the CPU so that it interacts with the video decoder front end processor

circuit 302 in such a manner as to produce fully decoded video data. Thus, under control of the routine 330 the CPU 316 performs video decoding operations not performed by the video processor circuit 302, e.g., motion
5 compensated prediction operations.

In the Fig. 2 embodiment, a display adapter 312 is used to couple the CPU 316 to the display device 210. The display adapter 312 converts video signals output
10 from the CPU 316 into a format used by the display device 210. The signals, e.g., video images, output by the display adapter are displayed on the display device 210.

Operation of the video processor circuit 302
15 will be described in detail below with regard to Fig. 4. However, generally speaking, the video processor circuit 302 receives configuration and other set-up information from the CPU 316 and processes encoded video received from the peripheral device 322 or CPU. The performed
20 processing includes, e.g., syntax parsing, variable length decoding, and/or other non-memory intensive video decoding operations. The video data, e.g., decoded intra-coded image data and/or partially decoded inter-coded image data, generated by video processor circuit
25 302, is passed to the CPU 316. In the Fig. 2 embodiment, the CPU 316 is responsible for completing the image decoding process to the extent that additional decoding operations are required. For example, in the case of inter-coded image data, the CPU 316 under control of the

video decoder routine 330 performs motion compensated predictions using reconstructed motion vector information received from the video processor circuit 302. As will be discussed below, in various embodiments, the video front end processor circuit 302 is capable of fully decoding intra-coded image data. Accordingly, in the case of intra-coded frames, the CPU 316 need not decode the image data output by the video front end processor circuit 302. However, when processing video data including both intra-coded and inter-coded frames, the previously decoded frame data output by the processor circuit 302 is normally stored on a temporary basis. The temporary storage of decoded image data allows for frame re-sequencing and renders the image data available in the event that it is needed as reference frame data in performing a motion compensated prediction operation.

Fig. 3 illustrates another exemplary embodiment of the present invention. The circuitry 400 illustrated in the Fig. 3 embodiment is similar to that of the Fig. 2 embodiment. However the display adapter 312 is replaced with a graphics processor 412 capable of performing and/or overseeing memory intensive video decoder operations. In addition, in the Fig. 3 embodiment a graphics memory 214 and a video bus 232 are included to facilitate performing memory intensive video operations such as motion compensated prediction. The graphics memory 214 is used for storing reference frame data used for motion compensated predictions while the video bus

232 is used for exchanging information between the main memory and video memory at a high rate of speed and without interfering with communications over the system bus 228.

5

Fig. 4 illustrates a video decoder front end processor circuit 302 which may be used in the systems illustrated in Figs. 2 and 3. As illustrated, the processor circuit 302 includes an input buffer 502, a VLD circuit 504, memory 520, an inverse scan circuit 506, inverse quantization circuit 508, inverse DCT circuit 510, an optional downsampler 512, a motion vector reconstruction circuit 518, a MUX 514, an output buffer 516, and an output control circuit 522.

15

While the circuit 302 includes memory 520, it is a relatively small amount of memory compared to the graphics memory 214 or main memory 326. This is because the memory 520 is used for storing control, configuration and, context information relevant to decoding operations as opposed to the actual frames. One kilobyte is an exemplary size of the memory 520 in the case of 8 bits per byte. The memory 520 includes a set of context, control, and configuration information 525, 527 for each of N bitstreams which are processed by the circuit 302. In addition to, or as part of, the memory 520, video buffers, e.g., processing buffers, may be included on the video decoder front end processor circuit 302.

25

The input buffer 502 is used to temporarily store video data and other input received by the video circuit 302 until the VLD circuit 504 is ready to process the input. The input can include encoded video data, configuration, and/or video bitstream information received from the peripheral device 322 and/or CPU 316. The encoded video data may be, e.g., an MPEG-2 packetized elementary stream (PES) or a video elementary stream which is produced by depacketizing a PES stream. Thus, the encoded video data input to the video circuit 302 may include, for example, coded motion vectors, coded quantized blocks of DCT coefficients, macroblock type information and other ancillary data such as coded block pattern information, coded macroblock position information, field select signals, and motion type information. Received configuration information may include, e.g., a reset signal, down sampling mode control information in the case of a system which supports downsampling, data output format (e.g., 8 or 9 bit output format information), and/or image region information indicating for example, whether a border region or a center region of a letterbox image is being sent to the circuit 302 for processing.

The VLD circuit 504 receives the input to the circuit 302 from the input buffer 502. The VLD circuit 504 is responsible for performing a variable length decoding operation on received variable length encoded data. The VLD circuit 504 also identifies ancillary

data. In addition, it identifies configuration, control, and video context information. Ancillary data and context information are supplied from the VLD circuit 504 to a third input of MUX 514. The configuration, control and video context information are also passed to the memory 520 and to the output control circuit 522.

Configuration, context, and/or control information can also be passed directly from the CPU 316 to the memory 520 thereby avoiding the need to transmit the information with the encoded video data being sent to the VLD circuit 504.

The memory 520 stores and outputs context information for each bitstream being decoded. The context information stored for each bitstream includes, e.g., high level syntax picture level and sequence level information. Exemplary stored picture level information includes frame type information (I, P, B frame information) while exemplary stored sequence level information includes vertical size, horizontal size, and frame rate. The video circuit 302 can be used to support decoding of multiple bitstreams, e.g., in order to support picture-in-picture capabilities. By storing context, control and configuration information 525, 527 for each of the multiple bitstreams being decoded the data need not be transmitted multiple times to the circuit 302. Each time a different bitstream is to be decoded, for which information is already stored in the

memory 520, stored information associated with the
bitstream to be decoded is supplied to the VLD circuit
504, motion vector reconstruction circuits 518 and
downsampling circuit 512. In this manner, the video
5 circuit 302 can be reconfigured to processes a different
bitstream in a fast and efficient manner. For example,
in response to the stored control information when
processing image data intended to be used as a main
picture, the downsampling circuit 512 may be disabled.
10 However, in response to stored control information when
processing inset picture data the downsampling circuit
512 may be enabled. Thus, data reduction operation may
be performed on inset picture data but not main picture
data in response to control information received from the
15 memory 520.

The memory 520 also supplies the control
information associated with a bitstream being processed
to the output control circuit 522. The output control
20 circuit 522 is coupled to a control input of the MUX 514
and a control input of the output buffer 516. The output
control circuit 522 controls which data, e.g., decoded
image data, motion vectors, ancillary data, etc. are
output. The circuit 522 also controls when the
25 information is output from the video circuit 302. The
output control circuit 522 does this by controlling the
MUX 514 and output buffer 516.

The VLD circuit 504 outputs image data generated by performing a variable length decoding operation. The image data output by the VLD circuit 504 is supplied to the inverse scan circuit 506 which
5 performs the inverse of the scan conversion process performed at encoding time. Next, the video data is processed by the inverse quantization circuit 508 and then the inverse DCT circuit 510. The inverse quantization circuit performs an inverse quantization
10 operation on received data while the inverse DCT circuit performs an inverse discrete cosine transform operation on received data. The decoded image data output by the IDCT circuit 510 is supplied to the downsampling circuit 512. Downsampling is performed in cases where the
15 received control and/or configuration information indicates that it should be applied. For example downsampling may be applied to inset picture data and/or HDTV images when the images will be displayed at lower than their full resolution. The resulting decoded image
20 data, corresponding to intra-coded frames prediction residuals, is supplied from the downsampler 512 to a first input of the MUX 514.

The VLD circuit 504 outputs motion vector
25 information including, e.g., motion vector values encoded using differential encoding, to the motion vector reconstruction circuit 518. The reconstruction circuit 518 reconstructs motion vectors from the received information, e.g., by performing a reverse differential

coding operation. The reconstructed motion vectors are supplied to a second input of the MUX 514.

As discussed above, the output control circuit
5 522 control the MUX 514. In response to the control
signal from the control circuit 522, the MUX 514 outputs
the data from one of its inputs. This information is
stored in the output buffer 516 prior to being
transmitted to the CPU 316 at a rate controlled by the
10 output control circuit 522. Thus, video processor
circuit 302 supplies the CPU 316 with decoded image data,
e.g., pixel values, corresponding to intra-coded frames
and intra-coded image portions. With regard to inter-
coded image portions it supplies, e.g., reconstructed
15 motion vectors, coded block pattern information or
pattern code information, macroblock position information
and/or pixel residual values used in performing motion
compensated predictions. In cases where the video
processor circuit 302 does not perform an IDCT operation,
20 it may also output to the CPU 316 DCT coefficients.

The CPU 316 or CPU 316 operating in conjunction
with a graphics processor 412 completes the decoding of
inter-coded images using the video data supplied thereto
25 from the video decoder front end processor circuit 302 of
the present invention.

In the above described manner, the video front
end processor circuit 302 of the present invention can be

used to enhance and/or increase the speed at which a system including a programmable processor can perform video decoding operations.

5 While the video decoder front end processor circuit 302 of the present invention has been described as performing several operations, it is to be understood that one or only a few of the described operations may be performed by a video front end processor with the
10 remaining operations being performed by a programmable CPU or a programmable CPU operating in combination with a graphics processor. The decoder front end processor may reside on a graphics processor board or on a separate printed circuit board, e.g., a computer plug in board.

15 Numerous additional embodiments, within the scope of the present invention, will be apparent to those of ordinary skill in the art in view of the above description and the claims which follow.